

Lecture 1: Course Introduction

Welcome

- This course is BIODS 220 Artificial Intelligence in Healthcare (CS 271, BIOMEDIN 220)
- What we hope you will get out of this course:
 1. Broad knowledge of opportunities for AI in healthcare
 2. Fluency in cutting edge deep learning algorithms, and practical ability to develop models for diverse types of healthcare data
 3. Understanding of real-world considerations and challenges for deploying AI algorithms in healthcare

Today's agenda

- A brief overview of AI in healthcare
- Course logistics

AI in healthcare: a rapidly exploding field




Hospitals Roll Out AI Systems to Keep Patients From Dying of Sepsis

Septic shock kills 50 percent of people who are affected—Sepsis Watch could save their lives

By **Eliza Strickland**


In hospitals, doctors and nurses keep vigilant watch over patients'



Google, Verily using AI to screen for diabetic retinopathy in India

The machine learning algorithm can also help with screening for diabetic macular edema, a boon for patients in a country where physicians are in short supply.

By **Mike Miliard** | February 26, 2019 | 03:17 PM



(Photo: Google)

AI in healthcare: a rapidly exploding field

Apple's future healthcare market moves will rely heavily on AI analysis

By Malcolm Owen
Monday, September 16, 2019, 09:03 am PT (12:03 pm ET)

Apple's moves in the healthcare market could involve the tracking of user data for further analysis by artificial intelligence and billing model based on cost-savings, with analysts pointing out areas of the consumer health industry Apple could easily advance by building upon its already-released technology and services.



Google to Store and Analyze Millions of Health Records

The tech company's deal with Ascension is part of a push to use artificial intelligence to aid health services.



SCIENCE BUSINESS TECH

Amazon is buying 'membership-based' healthcare provider One Medical for \$3.9 billion

19

One Medical's Netflex-for-primary-care is a \$199 subscription to a modern doctor's office

By Richard Lawler | @rjcc | Jul 21, 2022, 9:40am EDT | 19 comments

f t SHARE



SAMSUNG

Every angle is your best angle with FlexCam

LEARN MORE

MICROSOFT SCIENCE TECH

Microsoft Healthcare is a new effort to push doctors to the cloud


Microsoft wants to be a big part of the cloud and AI healthcare race

By Tom Warren | @tomwarren | Jun 27, 2018, 6:50am EDT

AI in healthcare: a rapidly exploding field

Google Trends Explore

● AI in healthcare
Search term

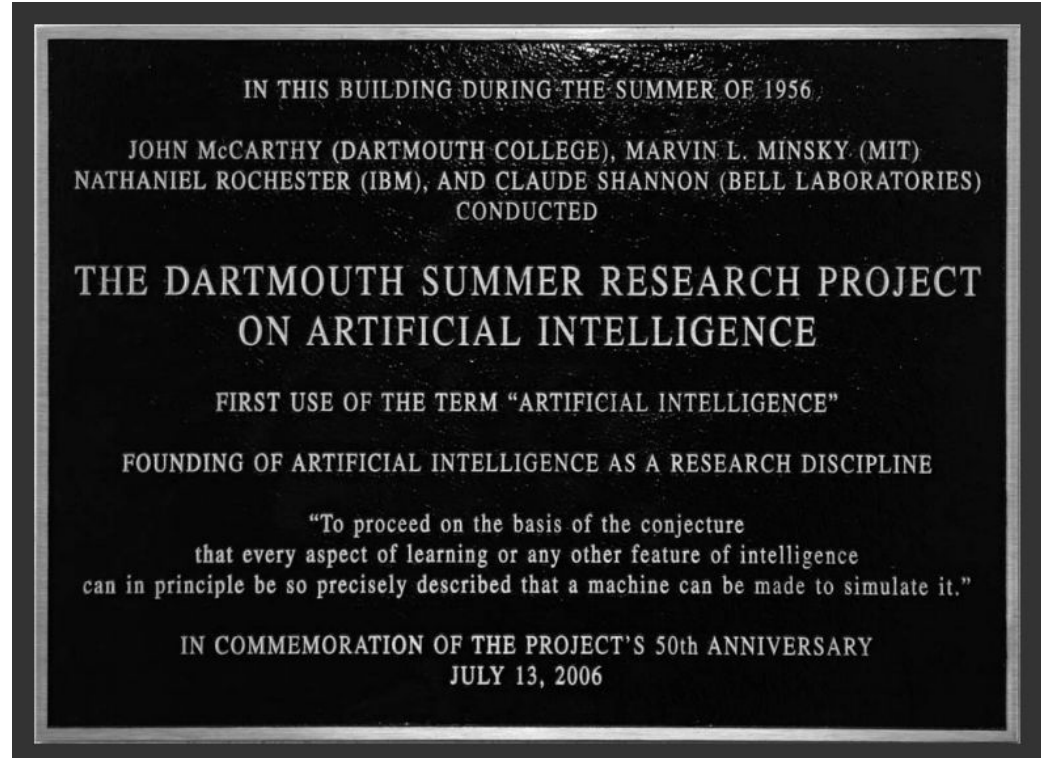


A journey back in time... brief history of modern AI



1956: Birth of AI as a modern research discipline



John McCarthy



Early progress in the late 50s and 60s



Perceptron model: Rosenblatt, 1958


```
Welcome to
          EEEEE LL   IIII ZZZZZZ  AAAAA
          EE    LL   II    ZZ   AA  AA
          EEEEE LL   II    ZZ   AAAAAA
          EE    LL   II    ZZ   AA  AA
          EEEEE LLLLL IIII ZZZZZZ  AA  AA

Eliza is a mock Rogerian psychotherapist.
The original program was described by Joseph Weizenbaum in 1966.
This implementation by Norbert Landsteiner 2005.

ELIZA: Is something troubling you ?
YOU:   Men are all alike.
ELIZA: What is the connection, do you suppose ?
YOU:   They're always bugging us about something or other.
ELIZA: Can you think of a specific example ?
YOU:   Well, my boyfriend made me come here.
ELIZA: Is it important to you that your boyfriend made you come here ?
YOU:   He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
YOU:   It's true. I am unhappy.
ELIZA: Can you explain what made you unhappy ?
YOU:   █
```

ELIZA chatbot: Weizenbaum, 1966

Progress and excitement in the late 50s and 60s




```
Welcome to
      EEEEE LL   IIII ZZZZZZ  AAAAA
      EE    LL   II    ZZ   AA  AA
      EEEEE LL   II    ZZ   AAAAAA
      EE    LL   II    ZZ   AA  AA
      EEEEE LLLLL IIII ZZZZZZ  AA  AA

Eliza is a mock Rogerian psychotherapist.
by Joseph Weizenbaum in 1966.
Steiner 2005.

suppose ?
something or other.

ELIZA: Can you think of a specific example ?
YOU: Well, my boyfriend made me come here.
ELIZA: Is it important to you that your boyfriend made you come here ?
YOU: He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
YOU: It's true. I am unhappy.
ELIZA: Can you explain what made you unhappy ?
YOU: █
```

Capabilities of early projects soon reached limitations they could not surpass.
Enthusiasm for AI dwindled in the 70s.



Perceptron model: Rosenblatt, 1958

ELIZA chatbot: Weizenbaum, 1966

Resurgence in the 80s

```
#DEFINE MOLFORM C 12 H 14 O
MOLECULAR FORMULA DEFINED
#DEFINE SUBSTRUCTURE Z
```

[Z is the structure required by constraints C2 and C3.]

■ ■ ■

CONSTRAINT: SUBSTRUCTURE CH0 EXACTLY 2

[C7: we must end up with exactly two quaternary carbons.]

CONSTRAINT: RING 3 NONE [C9]

CONSTRAINT: RING 4 NONE [C9]

CONSTRAINT:


#####

#####


47 STRUCTURES WERE OBTAINED

#DRAW ATNAMED (5 6)

[Comment: The following is a selection of final structures 5, 6.]



Expert systems, 1970s and 80s.
Feigenbaum, etc.



To get the correct generalization of the delta rule, we must set

$$\Delta_p w_{ji} \propto - \frac{\partial E_p}{\partial w_{ji}},$$

where E is the same sum-squared error function defined earlier. As in the standard delta rule it is again useful to see this derivative as resulting from the product of two parts: one part reflecting the change in error as a function of the change in the net input to the unit and one part representing the effect of changing a particular weight on the net input. Thus we can write


$$\frac{\partial E_p}{\partial w_{ji}} = \frac{\partial E_p}{\partial net_{pj}} \frac{\partial net_{pj}}{\partial w_{ji}}, \quad (9)$$

By Equation 7 we see that the second factor is

$$\frac{\partial net_{pj}}{\partial w_{ji}} = \frac{\partial}{\partial w_{ji}} \sum_k w_{jk} o_{pk} = o_{pi}. \quad (10)$$

Backpropagation. Rumelhart, 1986.

First appearances of modern neural networks



LeCun, 1990s.




Figure 1: Architecture of memory cell c_j (the box) and its gate units in_j, out_j . The self-recurrent connection (with weight 1.0) indicates feedback with a delay of 1 time step. It builds the basis of the "constant error carousel" CEC. The gate units open and close access to CEC. See text and appendix A.1 for details.

Schmidhuber, 1997.

First appearances of modern neural networks


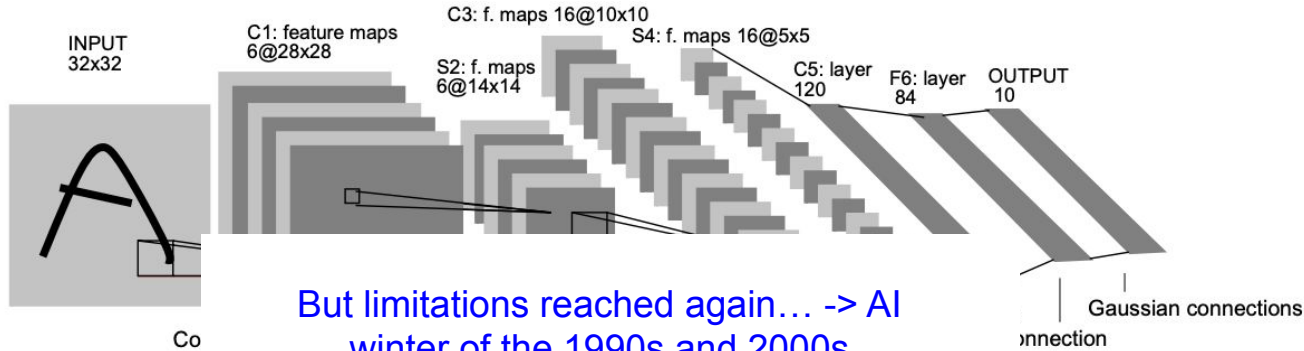

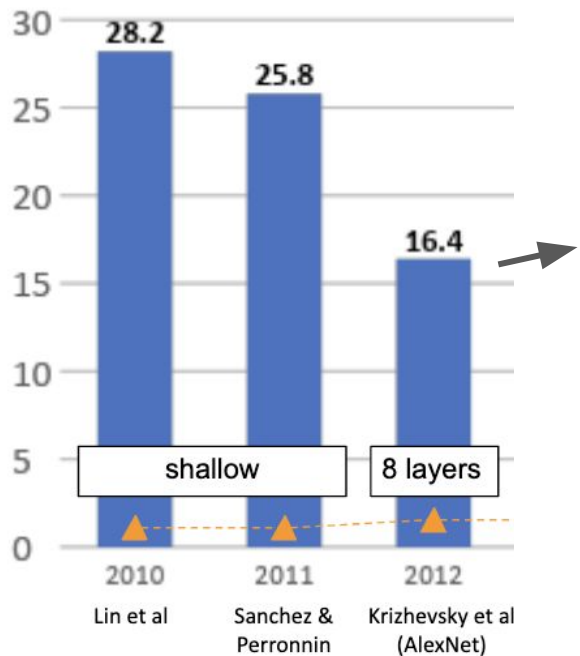


Figure 1: Architecture of memory cell c_j (the box) and its gate units in_j, out_j . The self-recurrent connection (with weight 1.0) indicates feedback with a delay of 1 time step. It builds the basis of the "constant error carousel" CEC. The gate units open and close access to CEC. See text and appendix A.1 for details.

Schmidhuber, 1997.

2012: Deep learning breakthrough




Krizhevsky et al. 2012. 8-layer “AlexNet”.


ImageNet Visual Recognition Challenge results.

Convergence of key ingredients of deep learning

Algorithms




Compute



Data



2015: Very deep convnets and challenging vision tasks

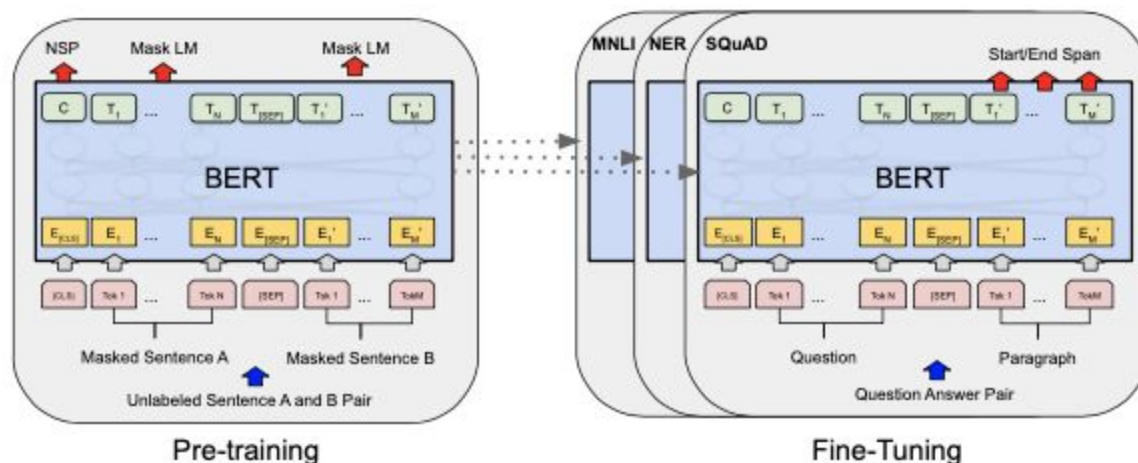


He et al. 2015. ResNet.



2018: Breakthroughs in deep learning for natural language processing (sequences)

Transformer architectures and pre-training -> fine-tuning. State-of-the-art on 11 NLP benchmarks.



Devlin et al. 2018. BERT.

2020: Very large scale text and image generation models

OpenAI models for text generation (left), text-to-image generation (right-top), and zero-shot classification tasks (right-bottom)

SYSTEM PROMPT (HUMAN-WRITTEN)

In a shocking finding, scientist discovered a herd of unicorns living in a remote, previously unexplored valley, in the Andes Mountains. Even more surprising to the researchers was the fact that the unicorns spoke perfect English.

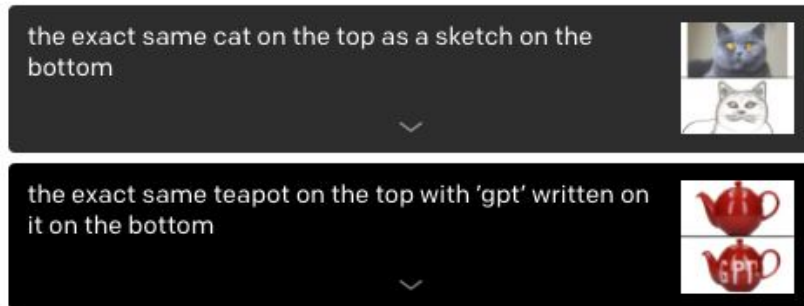
MODEL COMPLETION (MACHINE-WRITTEN, 10 TRIES)

The scientist named the population, after their distinctive horn, Ovid's Unicorn. These four-horned, silver-white unicorns were previously unknown to science.

Now, after almost two centuries, the mystery of what sparked this odd phenomenon is finally solved.


Dr. Jorge Pérez, an evolutionary biologist from the University of La Paz, and several companions, were exploring the Andes Mountains when they found a small valley, with no other animals or humans. Pérez noticed that the valley had what appeared to be a natural fountain, surrounded by two peaks of rock and silver snow.

GPT-3 (figure from GPT-2). Brown et al. 2020.



DALL-E. Ramesh et al. 2021.

healthy lymph node tissue (22.8%) Ranked 2 out of 2





this is a photo of lymph node tumor tissue

this is a photo of healthy lymph node tissue

CLIP. Radford et al. 2021.



Deep learning for healthcare: the rise of medical data



Deep learning for healthcare: the rise of medical data




X-rays (invented 1895).




Deep learning for healthcare: the rise of medical data



X-rays (invented 1895).




CT (invented 1972).




Deep learning for healthcare: the rise of medical data



X-rays (invented 1895).



CT (invented 1972).




MRI (invented 1977).

Deep learning for healthcare: the rise of medical data

Q: What are other examples of medical data?

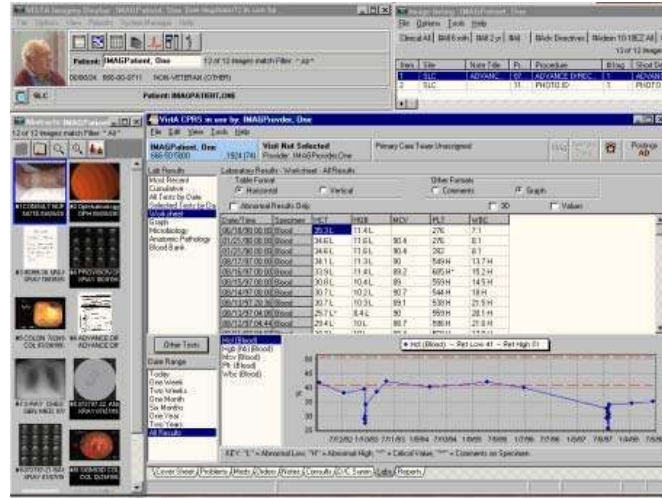
Electronic health records -- making patient data available



Imaging data



Patient measurements



1960s: invention
 1980s: increased effort
 2009: 51% adoption, HITECH Act
 2017: 98% adoption

Progress - CCC
 Note Date: 11/17/16
 Signed by (RHEUMATOLOGIST), MD on 11/21/16 at 11:00 am Affiliation: MEDICAL CENTER

Vital Signs sheet entries for 11/17/16: HR: 123/74, Heat Rate: 83, Weight: 173 (With Clothes), BMI: 26.9, Pain Score: 0

Active Medication list as of 11/17/16:


Medications - Prescription
 FLUCONAZOLE - fluconazole 0.01 % topical cream. Apply to affected area twice a day Use for up to 2 weeks as needed for flares.
 HYDROXYCHLOROQUINE - hydroxychloroquine 200 mg tablet. One tablet(s) by mouth daily
 INSULIN LEPRO HUMANLOGI - Humalog 100 unit/mL, subcutaneous cartridge. Insulin pump - (Prescribed by Other Provider)
 LEVOTHYROXINE - levothyroxine 75 mcg tablet. 1 tablet(s) by mouth qm
 LOSARTAN - losartan 50 mg tablet. 1 tablet(s) by mouth once a day am
 ROSUVASTATIN (CRESTOR) - Crestor 40 mg tablet. 1 tablet(s) by

Clinical notes


Test	Value	Reference
Hemoglobin	8.0	8.5–11.0 mmol/L
C-reactive protein	279	<5 mg/L
Red blood cell count	3.86	4.3–6.0 × 10 ¹² /L
White blood cell count	27.1	4.0–10.0 × 10 ⁹ /L
Thrombocytes	462	150–400 × 10 ⁹ /L
Glucose	12.9	4.0–7.8 mmol/L
Sodium	127	135–145 mmol/L
Potassium	4.2	3.5–5.0 mmol/L
Creatinine	40	50–110 μmol/L
Estimated glomerular filtration rate	>90	>60 ml/min
Ureum	3.2	2.5–7.5 mmol/L
Lactate dehydrogenase	166	<250 U/L
Aspartate aminotransferase	14	<40 U/L
Alanine aminotransferase	13	<50 U/L
Alkaline phosphatase	127	<120 U/L
Gamma-glutamyl transferase	96	<50 U/L

Lab results


Genomics data




1953 - Watson and Crick discover double helix structures of DNA




1977 - Fred Sanger sequences first full genome of a virus



1990 - 2003: Human Genome Project sequences full human genome



2003: ENCODE project launched to identify and characterize genes in human genome



2008 - 2015: 1000 Genomes Project International effort to study human genetic variation



2006 - present: UK Biobank Project Genetic data and intended 30 years of health follow-up for 500k individuals in the UK

Wearables and other sensor data



First iPhone: 2007





Fitbit: 2009





Apple Watch: 2014

AI in healthcare: biomedical image interpretation

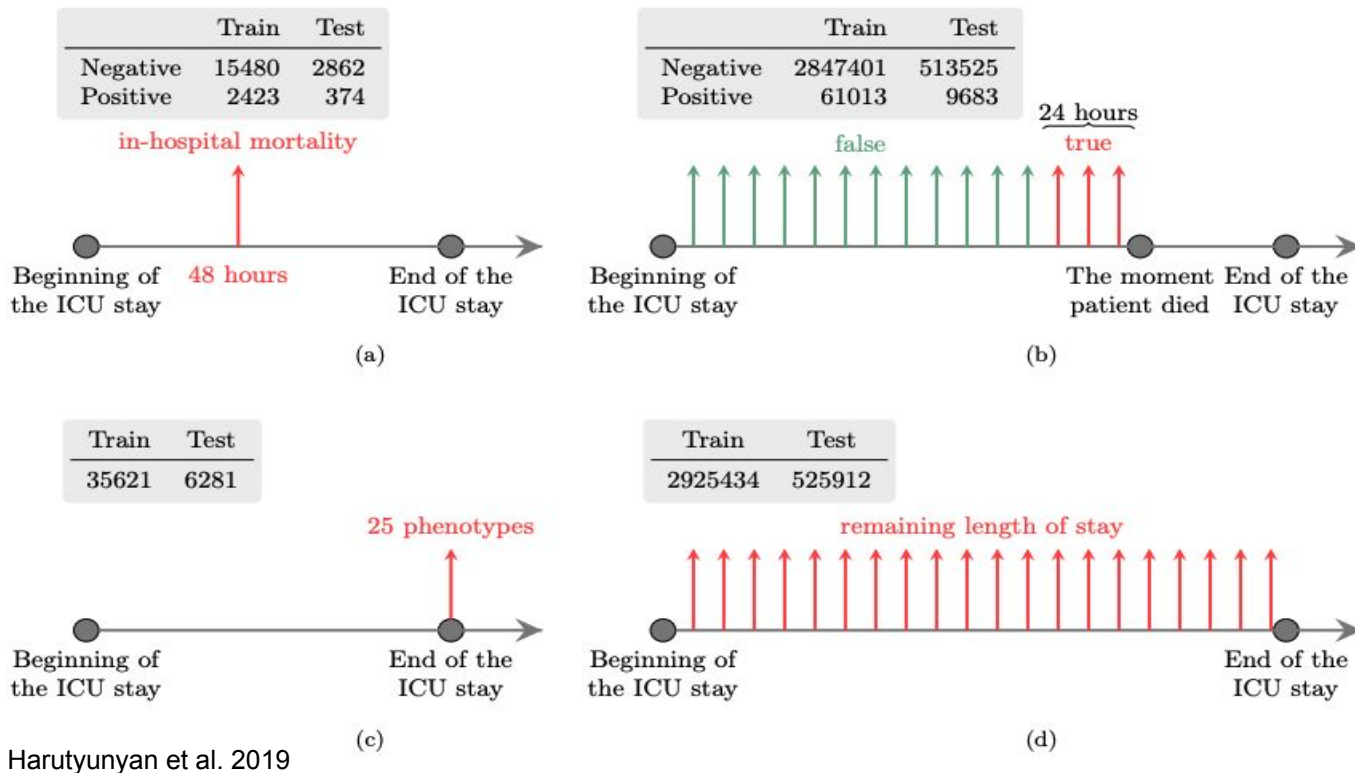


Wu et al. 2019

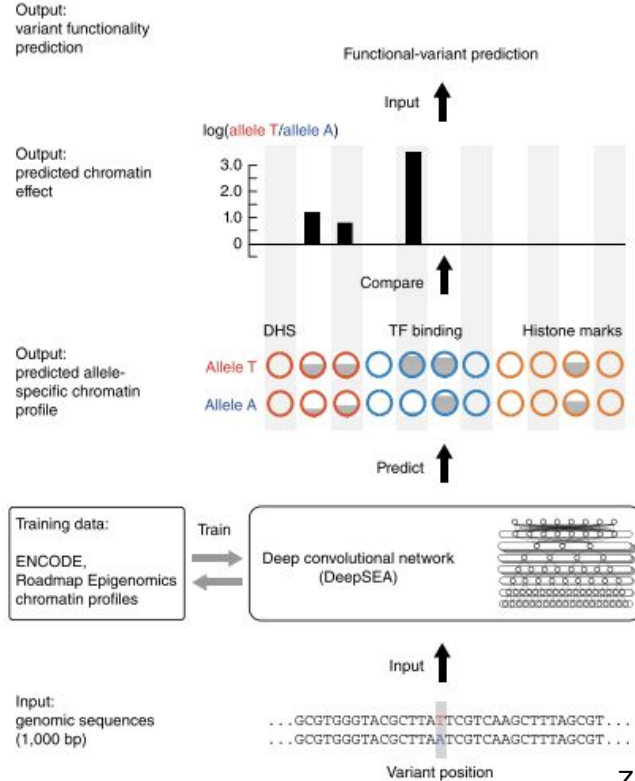


Liu et al. 2017


AI in healthcare: clinical event prediction



AI in healthcare: genomic analysis




Zhou et al. 2015



Poplin et al. 2016

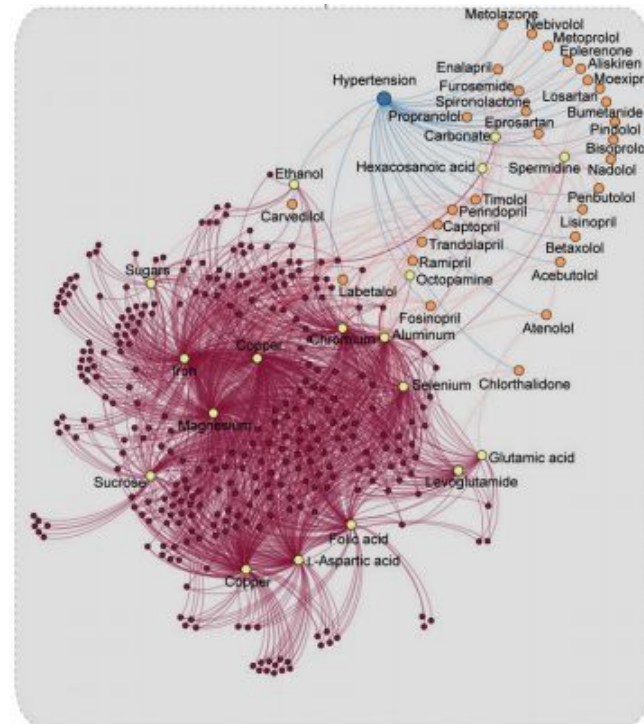
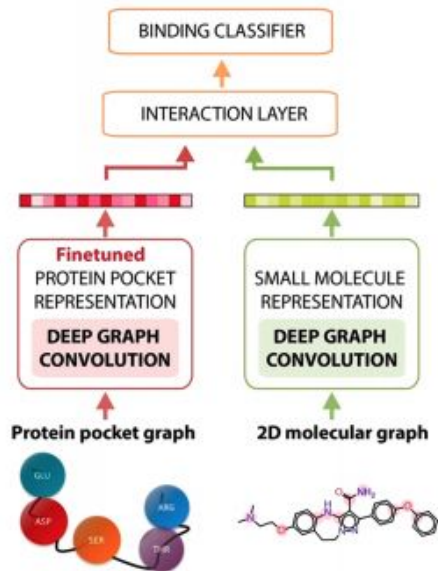
AI in healthcare: drug discovery and drug interaction prediction

Step I - Unsupervised Pocket Graph Autoencoder



Pocket Graph-CNN
Weight Initialization


Step II - Supervised Graph Convolutional Binding Classifier



Ryu et al. 2018

Tong et al. 2019

AI in healthcare: intelligent healthcare spaces and environments

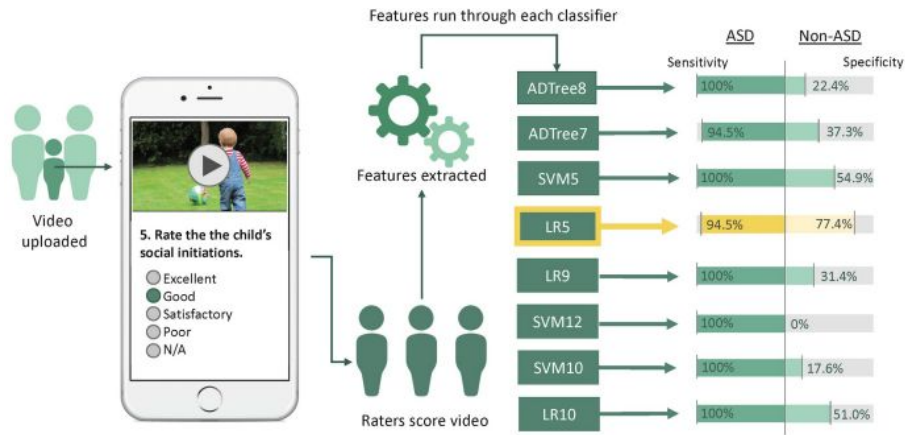


Yeung et al. 2019

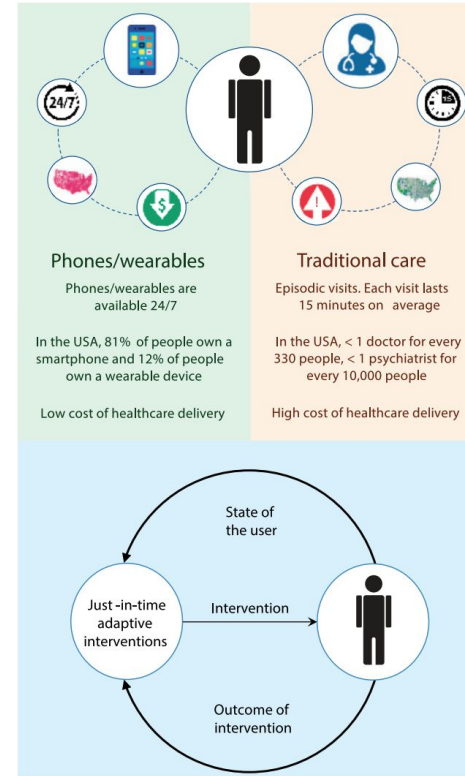


robinhealthcare.com

AI in healthcare: mobile health and wearables




Tariq et al. 2018




Menictas et al. 2019


AI in healthcare: recent applications for COVID-19



Harmon et al. 2020



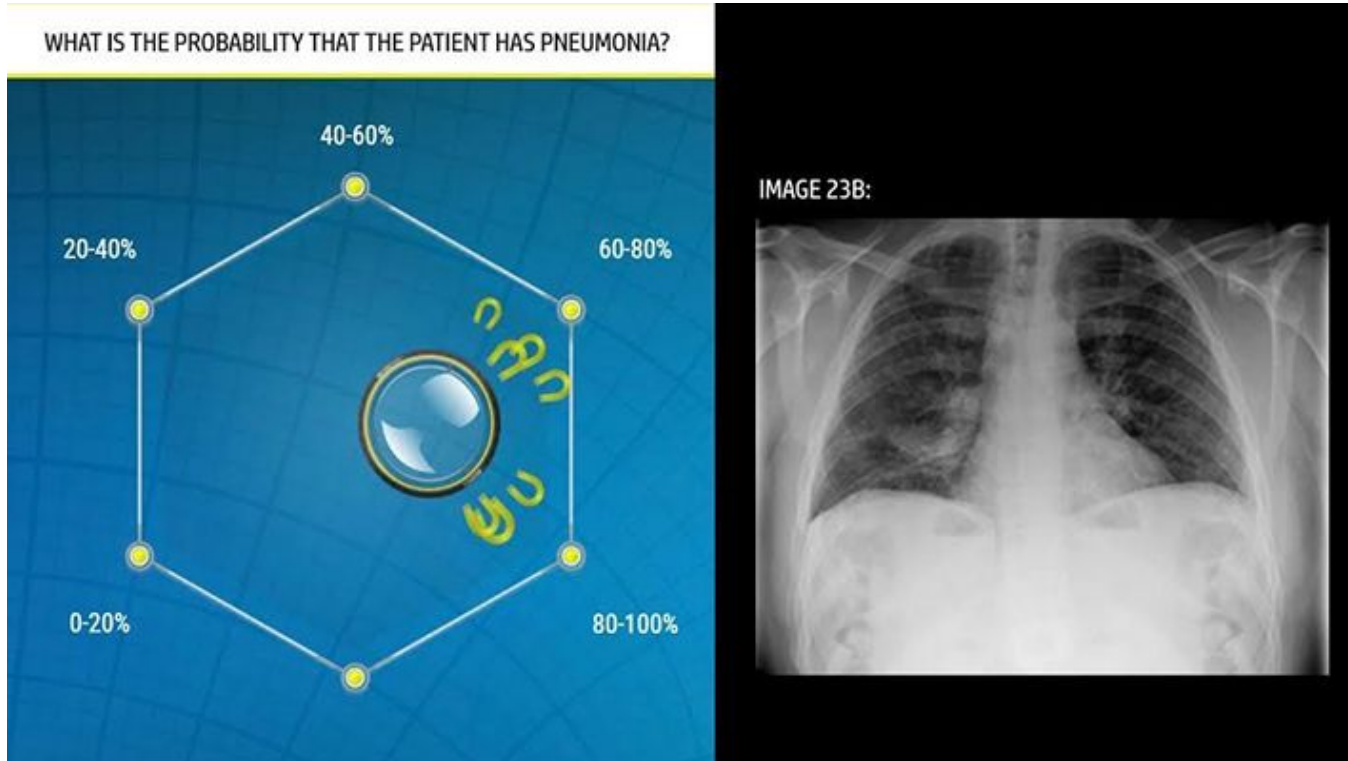
Yan et al. 2020



Jumper et al. 2020

The promise is great...
but many open challenges
in deployment as well

Uncertainty and AI / human collaboration



Rosenberg et al. 2018

Bias and fairness

RESEARCH

RESEARCH ARTICLE

ECONOMICS

Dissecting racial bias in an algorithm used to manage the health of populations

Ziad Obermeyer^{1,2*}, Brian Powers³, Christine Vogel⁴, Sendhil Mullainathan^{5*}†

Health systems rely on commercial prediction algorithms to identify and help patients with complex health needs. We show that a widely used algorithm, typical of this industry-wide approach and affecting millions of patients, exhibits significant racial bias: At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled illnesses. Remedying this disparity would increase the percentage of Black patients receiving additional help from 17.7 to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness, but unequal access to care means that we spend less money caring for Black patients than for White patients. Thus, despite health care cost appearing to be an effective proxy for health by some measures of predictive accuracy, large racial biases arise. We suggest that the choice of convenient, seemingly effective proxies for ground truth can be an important source of algorithmic bias in many contexts.

There is growing concern that algorithms may reproduce racial and gender disparities via the people building them or through the data used to train them (1–3). Empirical work is increasingly lending support to these concerns. For example, job search ads for highly paid positions are less likely to be presented to women (4), searches for distinctively Black-sounding names are more likely to trigger ads for arrest records (5), and image searches for professions such as CEO produce fewer images of women (6). Facial recognition systems increasingly used in law enforcement perform worse on recognizing faces of women and Black individuals (7, 8), and natural language processing algorithms encode language in gendered ways (9).

researcher-created algorithms (10–13). Without an algorithm's training data, objective function, and prediction methodology, we can only guess as to the actual mechanisms for the important algorithmic disparities that arise.

In this study, we exploit a rich dataset that provides insight into a live, scaled algorithm deployed nationwide today. It is one of the largest and most typical examples of a class of commercial risk-prediction tools that, by industry estimates, are applied to roughly 200 million people in the United States each year. Large health systems and payers rely on this algorithm to target patients for “high-risk care management” programs. These programs seek to improve the care of patients with complex health needs by providing additional


that rely on past data to build a predictor of future health care needs.

Our dataset describes one such typical algorithm. It contains both the algorithm's predictions as well as the data needed to understand its inner workings: that is, the underlying ingredients used to form the algorithm (data, objective function, etc.) and links to a rich set of outcome data. Because we have the inputs, outputs, and eventual outcomes, our data allow us a rare opportunity to quantify racial disparities in algorithms and isolate the mechanisms by which they arise. It should be emphasized that this algorithm is not unique. Rather, it is emblematic of a generalized approach to risk prediction in the health sector, widely adopted by a range of for- and non-profit medical centers and governmental agencies (21).

Our analysis has implications beyond what we learn about this particular algorithm. First, the specific problem solved by this algorithm has analogies in many other sectors: The predicted risk of some future outcome (in our case, health care needs) is widely used to target policy interventions under the assumption that the treatment effect is monotonic in that risk, and the methods used to build the algorithm are standard. Mechanisms of bias uncovered in this study likely operate elsewhere. Second, even beyond our particular finding, we hope that this exercise illustrates the importance, and the large opportunity, of studying algorithmic bias in health care, not just as a model system but also in its own right. By any standard—e.g., number of lives affected, life-and-death consequences of the decision—health is one of the most important and widespread social sectors in which algorithms are already used at scale today, unbeknownst to many.

Obermeyer et al. 2019

Privacy and security



Price et al. 2019



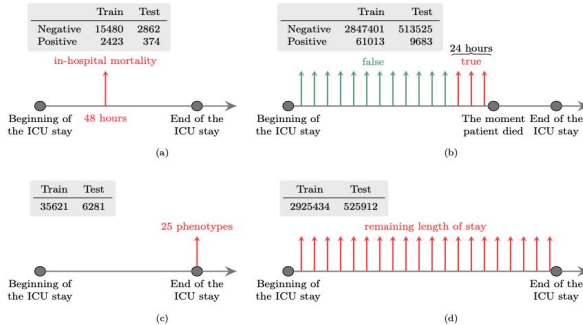



Figure: <https://news.developer.nvidia.com/first-privacy-preserving-federated-learning-system/>

In this class

1st part: developing DL algs for health data

2nd part: deploying AI for health



RESEARCH ARTICLE

ECONOMICS

Dissecting racial bias in an algorithm used to manage the health of populations

Ziad Obermeyer^{1,2*}, Brian Powers³, Christine Vogler^{4*}, Seshil Multanath^{1,2*}

Health systems rely on commercial prediction algorithms to identify and help patients with complex health needs. We show that a widely used algorithm, typical of this industry-wide approach and affecting millions of patients, exhibits significant racial bias. At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled diseases. Remedying this disparity would increase the percentage of Black patients receiving additional help from 17 to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness, but unequal access to care means that we spend less money caring for Black patients than for White patients. Thus, despite health care cost appearing to be an effective proxy for health by some measures of predictive accuracy, large racial biases arise. We suggest that the choice of correlated, seemingly effective proxies for ground truth can be an important source of algorithmic bias in many contexts.

There is growing concern that algorithms may reproduce racial and gender disparities via the people building them or through the data used to train them (2-5). Empirical work is increasingly finding support to these concerns. For example, job search ads for highly paid positions are less likely to be presented to women (6), searches for distasteful Black-sounding names are more likely to trigger ads for arrest records (5), and image searches for professions such as CEO produce fewer images of women (6). Facial recognition systems increasingly used in law enforcement perform worse for women (7, 8), and natural language processing systems increasingly used in customer service are less accurate for Black speakers (9).

Health systems rely on commercial prediction algorithms (10-15). With our algorithm's training data, objective function, and prediction methodology, we not only replicate as to the actual mechanisms for the important algorithmic disparities that arise. In this study, we exploit a rich dataset that provides insight into a live, scaled algorithm deployed nationwide today. It is one of the largest and most typical examples of a class of commercial risk-prediction tools that, by industry estimates, are applied to roughly 200 million people in the United States each year. Large health systems and carriers rely on

that rely on past data to forecast health care costs. Our dataset consists of 10 million patient records. It contains both the data used to train the model as well as the data used to evaluate the model's inner workings. It includes information used to fit objective functions, a set of outcome data, inputs, outputs, and data used to evaluate the model's performance. Our analysis has implications for the specific problem we have analyzed in this paper: the use of correlated proxies for ground truth in the development of risk prediction models. We believe that the insights from this study will be useful to other researchers and practitioners who are interested in the use of correlated proxies for ground truth in the development of risk prediction models.

Federated Server

Speed Breakouts

Get to know your classmates

2x 4-minute breakouts (4 students each)

- Name, program, year
- What's one thing you hope to get out of this class?
- What kind of healthcare tasks or data are you most interested in?

Course Logistics

Lectures: MW 1:30-2:50pm, Alway M106

- Lectures will be recorded and posted afterwards on Canvas


A few review sessions (e.g., Review of deep learning fundamentals, Tensorflow, Project partner finding, Midterm review): select Fridays 1:30-2:30pm, Alway M112

- First one will be a review of deep learning fundamentals, Fri 9/30
 - May be a little longer than usual, 1:30-2:50pm slot
- Stay tuned for announcements of future Friday reviews

Course materials will be hosted on website: <http://biods220.stanford.edu/>


Teaching team

Instructor



Serena Yeung
syyeung@stanford.edu
OH: Mon 9:30AM-11:30AM
Location: Packard 361

Teaching Assistants



Jeffrey Gu
jeffgu@stanford.edu
OH: Mon 3:00PM-5:00PM
Location: TBA



Ali Mottaghi
mottaghi@stanford.edu
OH: Thur 4:00PM-6:00PM
Location: TBA



Yuhui Zhang
yuhuiz@stanford.edu
OH: TBA
Location: TBA

Office hours will start week 2

Prerequisites

1. Proficiency in Python, or significant experience with a different programming language and ability to self-learn. Python will be used for homework assignments and the course project.
2. Basic familiarity with college calculus (e.g. Math 19 or 41, comfortable taking derivatives), linear algebra (e.g. Math 51 or EE 103 / CME 103, comfortable with common matrix vector operations and notation), and probability and statistics (e.g. CME 106 or CS 109, comfortable with common probability distributions, mean, standard deviation, etc).
3. Familiarity with machine learning, e.g. comfortable with the framework of machine learning and experience training a machine learning model.
4. Familiarity with deep learning is highly recommended, e.g. prior experience training a deep learning model. ****If you do not have this or do not feel comfortable with deep learning, this Friday's review session is critical.****

Ed Platform

- Will be used for **all** course communications.
- All students should be automatically added, please check
 - Periodically added from axess, may be a delay if you recently enrolled
 - Email staff mailing list biods220-aut2223-staff@lists.stanford.edu if you do not yet have access
- If it is a personal matter (e.g., OAE), please make a private post to the teaching team or instructor.
- Note: we will not be using canvas in this course, with the exception of posting course videos. Communications will be through Ed, and grading will be through gradescope.

Grading

- Sign up for Gradescope through the “logistics tab”
- Breakdown:
 - Assignment 1: 15%
 - Assignment 2: 15%
 - Assignment 3: 15%
 - Midterm (Mon Nov 7, in person): 15%
 - Course project: 40%

Assignments

- Main objective to build conceptual and practical foundations in using deep learning for biomedical data
- A0 (Data access prerequisites): Out Tue 9/27, due Tue 10/4. No grade, but required by due date to gain data access required for later assignments.
- A1 (Medical images): Out Tue 10/4, due Tue 10/18.
- A2 (EHR and text data): Out Tue 10/18, due Tue 11/1.
- A3 (Genomics): Out Tue 11/1, due Tue 11/15.

- In this class, all deadlines refer to 11:59pm PST on the stated day.
- A limited amount of Google cloud credits will be provided for the assignments. Should be sufficient, but use wisely.
- Collaboration policy: please read on course website. Study groups are allowed, but each student must produce independent assignment and write names of group on assignment.

Project

- Opportunity to gain in-depth experience developing an AI-based approach to a healthcare problem.
- Worth 40% of grade. Can work in groups of 1-3. (Grades will be calibrated by group size)
- Since large part of course is focused on deep learning, must involve implementation and training of at least one deep learning model on health data. Otherwise, significant flexibility in technical component (compare DL vs. non-DL models, analyze DL model in depth, novel DL architectures, etc.).
- Can use any health-related data of your choice. Options include public datasets and challenges (e.g., start from a published paper!), ongoing projects at Stanford (if applicable), projects suggestions from Stanford Medical School, etc.
- Will release detailed project guidelines and suggestions, and discuss in lecture next week.

Project (cont.)

- Graded components:
 - Proposal: Due Fri 10/21.
 - Milestone: Due Fri 11/18.
 - TA project advising sessions: after the milestone, details TBD.
 - Final project poster session: In person, during the final exam period for this course (Wed 12/14, 3:30-6:30pm)
 - Final report due: Fri 12/16.

Late days

- Can be used on A1, A2, A3, project proposal, project milestone report.
- Cannot be used on project final presentation, or final project report.
- 6 late days total, 2 max for any assignment.
- Grades will be deducted by 25% for each additional late day.

Course schedule

	Date	Topic
Lecture 1	Sep 26 (Mon)	Course Introduction
Assignment	Sep 27 (Tue)	Assignment 0 Released
Lecture 2	Sep 28 (Wed)	Medical Images: Classification
Section	Sep 30 (Fri)	Review: Deep Learning Fundamentals
Lecture 3	Oct 3 (Mon)	Medical Images: Advanced Vision Models (Detection and Segmentation)
Assignment	Oct 4 (Tue)	Assignment 0 Due; Assignment 1 Released
Lecture 4	Oct 5 (Wed)	Medical Images: Advanced Vision Models (3D and Video)
Section	Oct 7 (Fri)	Numpy/TensorFlow Review Session
Lecture 5	Oct 10 (Mon)	Electronic Health Records: Introduction
Lecture 6	Oct 12 (Wed)	Electronic Health Records: Advanced Topics
Section	Oct 14 (Fri)	Project Partner Finding Session (optional)
Lecture 7	Oct 17 (Mon)	Electronic Health Records: More on Text Data and Representations
Assignment	Oct 18 (Tue)	Assignment 1 Due; Assignment 2 Released
Guest Speaker	Oct 19 (Wed)	Strategies for Interdisciplinary Projects in AI and Healthcare
Project	Oct 21 (Fri)	Project Proposal Due
Lecture 8	Oct 24 (Mon)	Multimodal Data, Multimodal Models, Weakly and Self-Supervised Learning
Lecture 9	Oct 26 (Wed)	More on Transformers and Multimodal Models
Lecture 10	Oct 31 (Mon)	Genomics: Introduction
Assignment	Nov 1 (Tue)	Assignment 2 Due; Assignment 3 Released

Assignment	Nov 1 (Tue)	Assignment 2 Due; Assignment 3 Released
Guest Speaker	Nov 3 (Wed)	Genomics: Advanced Topics
Section	Nov 4 (Fri)	Midterm Review Session
Midterm	Nov 7 (Mon)	Midterm Exam (in person during class hours)
Guest Speaker	Nov 9 (Wed)	TBD
Lecture 11	Nov 14 (Mon)	Special Topics: AI for COVID-19
Assignment	Nov 15 (Tue)	Assignment 3 Due
Lecture 12	Nov 16 (Wed)	Unsupervised Learning and Reinforcement Learning
Project	Nov 18 (Fri)	Project Milestone Due
Lecture 13	Nov 28 (Mon)	Interpretability, Fairness, and Ethics
Lecture 14	Nov 30 (Wed)	Distributed Learning, Security, and Privacy
Guest Lecture	Dec 5 (Mon)	TBD
Project	Dec 14 (Wed)	Poster Session (in person from 3:30PM to 6:30PM in TBA)
Project	Dec 16 (Fri)	Project Report Due

Next time

- Start discussion of medical images